

TIME DIVISION MULTIPLEXING OF INTER-SYSTEM CHANNEL DATA STREAMS FOR TRANSMISSION ACROSS A NETWORK

Technical Field

[0001] The present invention relates in general to data communications, and more particularly, to a facility for time division multiplexing (TDM) at least two inter-system channel (ISC) data streams to create a common, TDM multiplexed data stream, while also preserving data disparity.

Background of the Invention

[0002] Fiber optic communication links are an integral part of many computer systems. Since it is quite expensive to construct new fiber optic lines, various approaches have been devised to increase system capacity. These approaches include optical wavelength division multiplexing (WDM), dense wavelength division multiplexing (DWDM), coarse wavelength division multiplexing (CWDM) and, for certain data formats, time division multiplexing (TDM).

[0003] In a WDM or DWDM system, multiple optical channels are carried over a single optical fiber, each channel being assigned a particular optical wavelength. By using a single, existing optical fiber to carry multiple channels, DWDM has been found to be an efficient approach for increasing the capacity of data communication systems.

[0004] Each channel in a current DWDM system can carry data at speeds up to 10 Gbit/s. This works well for data that is being transmitted at such speeds, however, it is obviously inefficient to use a channel capable of 10 Gbit/s to transmit a signal at a substantially lower data rate.

[0005] TDM modules have been employed, for certain data formats, in an attempt to avoid such inefficiency. At the transmitting end of a TDM system, an electronic switch (multiplexing unit) picks up signals from multiple input channels, in a predetermined

“channel-by-channel” order. A resulting multiplexed signal, which combines such input signals, is constructed and forms the “output” or transmitting end of the TDM system. The multiplexed signal is distributed to receiving terminal equipment as the output of the transmitting end of the module. TDM is discussed in detail in various publications, including a textbook by A.B. Carlson and D.G. Gisser entitled: Electrical Engineering Concepts and Applications, published by Addison-Wesley, N.Y. (1991), which is hereby incorporated herein by reference in its entirety.

[0006] Existing TDM modules typically work with a specific fixed data format. For example, modules have been designed to operate with the FDDI data format, while others work only in the Ethernet data format.

Summary of the Invention

[0007] Presented herein is an approach to extending the time division multiplexing (TDM) application to inter-system channel (ISC) links. Further, the extension is accomplished in such a way so as to preserve data disparity.

[0008] Provided, in one aspect, is a method which includes: time division multiplexing (TDM) at least two inter-system channel (ISC) data streams to create a TDM multiplexed data stream; and outputting the TDM multiplexed data stream for forwarding across a network.

[0009] In further aspects, the network may comprise a wavelength division multiplexing (WDM) network, and the method may further include wavelength division multiplexing the TDM multiplexed data stream with at least one other data stream onto a single wavelength for forwarding across the WDM network. Still further, the time division multiplexing can employ a first-in first-out (FIFO) buffer, and include at least one of inserting and deleting words in at least one data stream of the at least two ISC data streams to balance input and output FIFO buffer clock rates. This at least one of inserting and deleting of words is performed so as to maintain disparity balance within the least two ISC data streams. In addition, the at least two ISC data streams may include at least

some of idle sequences, frame sequences and continuous sequences, and the at least one of inserting and deleting can include employing a first protocol for the idle sequences and frame sequences, and a second protocol for the continuous sequences within the ISC data stream. Examples of the first and second protocol are described and claimed.

[0010] Systems and computer program products corresponding to the above-summarized methods are also described and claimed herein.

[0011] Additional features and advantages are realized through the techniques of the present invention. Other embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed invention.

Brief Description of the Drawings

[0012] The subject matter which is regarded as the invention is particularly pointed out and distinctly claimed in the claims at the conclusion of the specification. The foregoing and other objects, features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

[0013] FIG. 1 depicts one example of a networked computer environment to employ multiplex/demultiplex systems in accordance with an aspect of the present invention;

[0014] FIG. 2 depicts one partial embodiment of a networked computing environment employing a wavelength division multiplexing (WDM) multiplex/demultiplex system for inter-system channel (ISC) links;

[0015] FIG. 3 depicts one partial embodiment of a networked computing environment utilizing a time division multiplexing (TDM) multiplex/demultiplex system for ISC links, in accordance with an aspect of the present invention;

[0016] FIG. 4 depicts one partial embodiment of a networked computing environment employing a WDM multiplex/demultiplex system and a TDM multiplex/demultiplex system for ISC links, in accordance with an aspect of the present invention;

[0017] FIG. 5 is a flowchart of one embodiment of processing implemented by a TDM multiplex/demultiplex system for evaluating need for and performing frequency compensation within idle sequences or frame sequences in one or more of at least two ISC data streams being time division multiplexed, in accordance with an aspect of the present invention; and

[0018] FIG. 6 is a flowchart of one embodiment of processing implemented by a TDM multiplex/demultiplex system for determining need for and performing frequency compensation within continuous sequences in one or more of at least two ISC data streams being time division multiplexed, in accordance with an aspect of the present invention.

Best Mode for Carrying Out the Invention

[0019] Dense wavelength division multiplexing (DWDM) has emerged as a principal approach for transmitting tens to hundreds of independent data links over only a few optical fibers. This approach is employed by International Business Machines Corporation in metropolitan area networks (50-100 km distances) for disaster recovery applications, such as the Geographically Dispersed Parallel Sysplex (GDPS) application.

[0020] IBM's Parallel Sysplex architecture, such as the z/900 and z/800 Enterprise Servers, employs coupling links referred to as inter-system channel (ISC) Parallel Sysplex coupling links. A data stream passing on one of these links is referred to herein as an ISC data stream. Various aspects of ISC links have been described in numerous prior publications, including United States Letters Patent Nos. 6,438,285; 6,359,713;

6,359,709; and 6,356,367, each of which is hereby incorporated herein by reference in its entirety.

[0021] Currently available DWDM solutions are capable of multiplexing up to 32 wavelengths over a single fiber pair. Recently, the data rate of a single wavelength has increased into the 2.5 – 10 Gbits/s range, making it possible to time division multiplex (TDM) multiple data channels over a single wavelength. Thus, presented herein is an approach for time division multiplexing at least two data streams in inter-system channel (ISC) data format, for example, for use in a GDPS. Advantageously, significant cost savings can be realized by the techniques presented herein, for example, by allowing multiple ISC channels to share a single wavelength in a DWDM network.

[0022] FIG. 1 depicts one example of a networked computer environment, generally denoted 100 to employ multiplex/demultiplex systems 130, 150, which can be implemented in accordance with aspects of the present invention. Computer environment 100 includes a local server 110, which employs multiple links 120 in sending and receiving data across a network 140. Multiplex/demultiplex system 130, such as a wavelength division multiplexing system, is employed to increase transmission capacity of optical fibers comprising part of network 140. Similarly, multiplex/demultiplex system 150 is employed between network 140 and a remote server 170, which receives and transmits data across multiple links 160. As noted, in one implementation, networked computer environment 100 may employ inter-system channel (ISC) links, for example, as links 120 & 160 from/to local server 110 and remote server 170, respectively. ISC coupling links are marketed, for example, by International Business Machines Corporation for use in communicating between senders and receivers, such as within a Parallel Sysplex environment.

[0023] FIG. 2 depicts one implementation for a multiplex/demultiplex system facility 200 interfacing between multiple ISC links and a network 210. This multiplex/demultiplex system comprises a wavelength division multiplexing (WDM) system wherein a plurality of ISC links or channels are multiplexed onto a WDM link for

transmission of data across network 210. For example, currently 32 wavelengths may be mapped onto to a single physical cable, with each wavelength being dedicated to a particular ISC link.

[0024] FIG. 3 depicts a partial embodiment of a network computing environment wherein a time division multiplexing (TDM) multiplex/demultiplex system 300 is employed to interface between at least two ISC links and a network 310. There are a number of considerations when attempting to time division multiplex ISC channel protocols. For example, a TDM implementation commonly uses a first-in, first-out (FIFO) buffering scheme. Time division multiplexing will break down if the outbound data clocking rate differs from the inbound clock rate. One way to compensate for differences is to either insert or delete words into one or more of the ISC data streams. In order to accomplish this without violating data disparity, it is necessary to account for different data structures on an ISC channel and the effects the scheme will have on the attached processors.

[0025] Before describing the ISC data stream protocol in greater detail, FIG. 4 depicts an alternate embodiment of a network computer environment wherein two or more ISC data streams are TDM multiplexed by system 400 to create a TDM multiplexed data stream 405, which is then wavelength division multiplexed by WDM multiplex/demultiplex system 410 with one or more additional data streams 415 before forwarding across a network 420. Additional data streams 415 could themselves comprise TDM multiplexed data streams, with each TDM multiplexed data stream being assigned a particular wavelength of the available wavelengths of a WDM (or DWDM or CWDM) link. By way of specific example, TDM multiplex/demultiplex system 400 could implement a 2:1 TDM at 2.5 Gbits/s or an 8:1 TDM at 10 Gbits/s, while WDM multiplex/demultiplex system 410 may implement a 32:1 wavelength division multiplexing of data streams onto a TDM multiplexed link.

[0026] Returning to the ISC protocol, today's ISC channel comprises 1024 words, 4 bytes each, for a total of over 4000 bytes (plus some overhead) in a frame. This includes

a 4 byte idle word, 4 bytes of data, a 4 byte CRC check, 4096 bytes of data, another 4 byte CRC check, and a 4 byte idle word. ISC channels were originally designed to adjust their data transfer rate to match the processor cycle time by inserting null words into the data stream. An ISC-1 channel link needs to maintain at least 4 idle words between successive data frames. It should always be safe to insert null words in the data stream, but as a safeguard to preserve parity this should be done in pairs. Null words should be deleted in pairs for the same reasons. It should be noted that some non-IBM implementations of ISC channels will check for valid disparity of the null words, so this part of the architecture should be maintained.

[0027] In order to maintain a minimum safe idle spacing between data frames, the TDM system should wait for a sequence of at least 8 continuous idle words and delete one pair to leave a minimum spacing of 6 words (the minimum of 4 plus a margin of 2). This should be possible for most links which are at low utilization, though care must be taken to not run out of idles if they are being deleted from multiple points along the networked computer environment. This imposes a restriction that the FIFO overrun prevention must not require more idles to be inserted than the maximum number of add/drop nodes in a network (for example, 2 nodes in a point-to-point network).

[0028] There are three types of data transfer which can occur in the ISC channel: idles, frames, and continuous sequences. For the first two types, the approach above-described of inserting/deleting two null words is sufficient to maintain disparity. In the final case, it may be necessary to delete both a pair of idles and a data word in order to keep the proper disparity. The idle and null format permits them to be used in this way.

[0029] One method for frequency tolerance compensation (in accordance with an aspect of the present invention) is outlined below using the following notation:

/I/ = an idle word
/N/ = a null word
/D/ = a dataword within a Frame
/C/ = a dataword within a Continuous sequence

/Cx/ = an explicit dataword [C1 through C4] corresponding to the 4 types of dataword used in Continuous sequences

(i) Compensation within Idle Sequences (inter-frame-gap)

Insertion:

Wait for a sequence of two */I/* words, ignoring any */N/* words,
Add two */I/* words.

Deletion:

Wait for a sequence of eight */I/* words, ignoring any */N/*
Remove a desired number of */N/* words which occur between 2 consecutive */I/* words being removed,
Remove in pairs, a desired number of */I/* words.

(ii) Compensation within Continuous Sequences

Insertion:

Wait for one sequence of */C/I/C/I/*, ignoring any */N/* words
Add a */C/I/* word-pair (replicate the last */C/I/* pair).

Deletion:

Wait (ignoring any */N/*) for a sequence of */Cx/I/Cx/I/*
Remove one */Cx/I/* word-pair.

[0030] For (i) Deletion described above, this approach will remove any */N/* word which might intervene inbetween the 2 */I/* words which are being removed. This is significantly easier to implement than preserving the */N/* word in the data stream. For example, consider the following:

Example sequence before idle deletion:
/D/D/D/D/I/N/I/I/N/I/I/I/I/N/N/N/I/

Example sequence after idle deletion:
/D/D/D/D/I/N/I/I/N/I/I/I/

[0031] The (ii) Deletion approach is acceptable, even though it can potentially leave just one */Cx/I/* pair in the stream if there were only 2 such pairs in the stream to start with.

Example before deletion:
/D/D/D/D/I/N/I/N/C2/I/C2/I/N/N/N/I

After idle deletion:
/D/D/D/D/I/N/I/N/C2/I/I/N/N/N/I/

[0032] FIGS. 5 & 6 depict flowchart examples for frequency compensation within idle sequences or frames, as well as continuous sequences, respectfully. Beginning with FIG. 5, frequency compensation within idle sequences or frames 500 begins with a TDM system checking the ISC data streams' input data rates 510. In this discussion, an assumption is made that the TDM system is continuously monitoring buffer inputs so that the particular sequence within the data stream is known at any given time. Thus, the particular frequency compensation of FIG. 5 or FIG. 6 is readily selected as needed. The TDM system determines whether it is necessary to insert or delete words within a sequence 520, and if not, the time division multiplexing function proceeds 530. Otherwise, processing determines whether words are to be deleted 540. If no, then two idle words 550 are added. Otherwise, processing waits for a sequence of eight idle words, ignoring any null words 560. Once such a sequence is identified, a desired number of null words occurring between two consecutive idle words are removed 570. Null words are removed between two idle words which are also to be removed. Idle words are removed in pairs 580.

[0033] In FIG. 6, frequency compensation within a continuous sequence 600 begins with the TDM system checking the ISC data streams' input data rates 610. Processing determines whether inserting or deleting of words is required 620, and if not, continues with TDM function processing 630. Otherwise, the TDM system waits for a repeated sequence of continuous word idle word pairs, ignoring any null words 640. Once detected, processing determines whether words are being added or deleted 650. If added, then a continuous word idle word pair is added, replicating the last such pair 660. Otherwise, processing waits for a repeated sequence of continuous word idle word pairs 670, and removes one such pair 680.

[0034] Latency is another consideration for TDM multiplexing of ISC channels. The use of TDM multiplexers introduces additional latency on the data links, which depends on the design of the time division multiplexer, the data rate of the signals being multiplexed, and the frame sizes being multiplexed. For faster processors, for example, processors in the IBM z/900 server line, latency has a significant impact on system performance. For 2 Gbits/s ISC links, low latency is more important to system performance than the increased bandwidth available compared with 1 Gbits/s ISC links. It is estimated that time division multiplexing of ISC channels through a pair of DWDM devices in a point-to-point network will incur an additional latency which should be acceptable without impacting performance compared with an approach which does not time division multiplex the ISC channels.

[0035] The present invention can be included in an article of manufacture (e.g., one or more computer program products) having, for instance, computer usable media. The media has embodied therein, for instance, computer readable program code means for providing and facilitating the capabilities of the present invention. The article of manufacture can be included as a part of a computer system or sold separately.

[0036] Additionally, at least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform the capabilities of the present invention can be provided.

[0037] The flow diagrams depicted herein are just examples. There may be many variations to these diagrams or the steps (or operations) described therein without departing from the spirit of the invention. For instance, the steps may be performed in a differing order, or steps may be added, deleted or modified. All of these variations are considered a part of the claimed invention.

[0038] Although preferred embodiments have been depicted and described in detail herein, it will be apparent to those skilled in the relevant art that various modifications, additions, substitutions and the like can be made without departing from the spirit of the

invention and these are therefore considered to be within the scope of the invention as defined in the following claims.